arXiv:1605.05659v3 [math.DS] 2 Feb 2017

CUTTING SEQUENCES ON SQUARE-TILED SURFACES

CHARLES C. JOHNSON

Abstract. We characterize cutting sequences of infinite geodesics on square-tiled surfaces by considering interval exchanges on specially chosen intervals on the surface. These interval exchanges can be thought of as skew products over a rotation, and we convert cutting sequences to symbolic trajectories of these interval exchanges to show that special types of *combinatorial lifts* of Sturmian sequences completely describe all cutting sequences on a square-tiled surface. Our results extend the list of families of surfaces where cutting sequences are understood to a dense subset of the moduli space of all translation surfaces.

1. Introduction

Given a surface with a Riemannian metric and a labeled collection of curves on the surface, the *cutting sequence* of a geodesic is a biinfinite sequence of labels given in the order in which the geodesic intersects the corresponding curves. The most obvious question to consider concerns characterizing precisely which sequences of labels correspond to cutting sequences.

This question has been studied for a few special types of translation surfaces, which are surfaces equipped with a flat metric with cone-type singularities and trivial holonomy. In particular, cutting sequences have been described for flat tori [MH38] [Ser85]; the regular polygon surfaces with an even number of sides were studied by Smillie and Ulcigrai, [SU11]; Davis later extended the methods of Smillie and Ulcigrai to double n-gon surfaces with odd n [Dav13]; the L-shaped surface built from three squares was studied by Wu and Zhong, [WZ15]; and most recently the family of Bouw-Möller surfaces was studied by Davis, Pasquinelli, and Ulcigrai [DPU15]. Most of the previous results are obtained by considering the action of an element of the Veech group (group of affine symmetries) of the surface. In [Dav14], for example, the effect of *flip-shears* on cutting sequences on some special types of surfaces are considered. In this paper we add one more family of surfaces to this list by characterizing cutting sequences on square-tiled surfaces.

Our approach will be to convert the question of which sequences are cutting sequences into a question about the languages determined by a special class of interval exchange transformations which can be described as a certain type of skew product over a rotation. Languages of the family of interval exchanges with a fixed permutation were studied in [FZ08].

Date: February 3, 2017.

1991 *Mathematics Subject Classification.* Primary 37E35.

Key words and phrases. Cutting sequences, dynamical systems, square-tiled surfaces, and translation surfaces.

The author thanks the referee for their careful reading of the original manuscript and for their many helpful remarks for improving the exposition.

After recalling the pertinent background and establishing some vocabulary and notation we will prove that all cutting sequences on a given square-tiled surface can be described as sequences of pairs with the first entry in each pair giving the label of some square on the surface, and the second entry giving an edge of the square. The second entries of these pairs must form a cutting sequence on a torus, and the first entries must form a sequence of labels that is consistent with the way the various squares on the surface are glued together. The sequences of pairs satisfying these requirements have a relatively simple combinatorial description, but not all such sequences are cutting sequences on the surface as this collection includes sequences which would describe a geodesic which hits a conical singularity on the surface. We will show there is a combinatorial description of these sequences by considering walks on a family of graphs associated to the surface.

We describe the necessary background on square tiled surfaces, Sturmian sequences, languages, and interval exchanges in Section 2. We then show in Section 3 that interval exchanges for some specially chosen intervals on a square-tiled surface are simply skew products over rotations which greatly simplifies understanding these interval exchanges. In Section 4 we introduce the notion of consistency between sequences of edges and the gluings of the surface and describe symmetric Sturmian sequences. The complete statement and proof of the characterization theorem appears in Section 5.

2. Background

2.1. **Translation surfaces and the Veech dichotomy.** A *translation surface* is a closed, oriented surface equipped with a geometry which is flat away from a discrete set of *conical singularities* where small neighborhoods of a singularity are isometric to a Euclidean cone of cone angle $2\pi n$ with $n \geq 1$ an integer. Such a surface can always be represented as a collection of polygons in the plane together with gluings that identify the edges of polygons in pairs by translations where the inward-pointing normals of two identified edges point in opposite directions. Equivalently, a translation surface is a pair (X, ω) where X is a compact Riemann surface and ω is a holomorphic 1-form on X. The equivalence between these two definitions comes from endowing (X, ω) with a singular flat metric by integrating ω to obtain an atlas of charts where transitions are given by translations (hence the name *translation surface*), and so the usual Euclidean metric of $\mathbb{C}$ can be pulled back to a metric on the surface. Choosing a geodesic triangulation of the surface where all singularities occur as vertices of the triangulation then gives a description of (X, ω) as a collection of polygons with edge gluings.

The set of all translation surfaces where the 1-form ω has the orders of its zeros fixed is called a *stratum* and is equipped with the structure of a complex orbifold by assigning local coordinates by integrating ω along the relative cycles of each surface's first homology $H_1(X, \{p_1, p_2, ..., p_n\}; \mathbb{C})$, relative with respect to the zeros of the 1-form. The group $\mathrm{GL}(2, \mathbb{R})$ acts on each stratum by deforming the polygonal representation of a surface. Since these deformations are linear, parallel lines are preserved and so the deformed polygon still glues together to give a translation surface in the same stratum.

The stabilizer of a surface (X, ω) under the $\mathrm{GL}(2, \mathbb{R})$ action is called the *Veech group* of the surface and is denoted $\mathrm{SL}(X, \omega)$. It was shown in [Vee89] that when $\mathrm{SL}(X, \omega)$ is a lattice, then for every direction the geodesic flow on (X, ω) satisfies

the following dichotomy: the flow is either completely periodic or uniquely ergodic. It had been previously shown in [KMS86] that for every translation surface the flow in almost every direction is uniquely ergodic, but there are surfaces where the flow in some directions is minimal but not uniquely ergodic. *Veech surfaces*, the surfaces with a lattice stabilizer, are surfaces where this is impossible: every non-periodic direction is uniquely ergodic.

More information about translation surfaces can be found in the surveys [MT02], [Zor06], and [Wri15].

2.2. **Square-tiled surfaces.** Let $\mathbb{T}^2$ denote the standard square torus, $\mathbb{C}/\mathbb{Z}[i]$ equipped with the 1-form dz. Notice that the torus is a Veech surface with Veech group $\mathrm{SL}(2,\mathbb{Z})$. A *square-tiled surface* is a covering $\pi : X \to \mathbb{T}^2$ branched over at most one point. The cover X is then given the translation structure corresponding to $\omega = \pi^*(dz)$. We can describe X as a collection of unit squares, the number of squares being the degree of the covering, glued together so that the top of each square is identified with the bottom of another (or the same) square, and similarly for left- and right-hand edges of the squares. The gluings of the squares are determined by the covering's monodromy.

Every square-tiled surface determines, and is determined by, a pair of permutations. Supposing the surface is made of d squares, let $\Lambda = \{1, 2, ..., d\}$ be the labels of the squares. We then consider two permutations, h and v, of Λ and consider the surface obtained by gluing the right-hand edge of square λ to the left-hand edge of square $h(\lambda)$, and the top of square λ to the bottom of square $v(\lambda)$. This surface will be connected if and only if the group $\langle h, v\rangle$ acts transitively on Λ, and we will only consider connected surfaces.

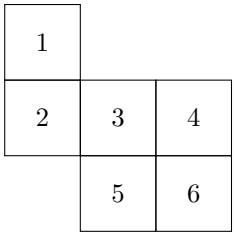

FIGURE 1. The square-tiled surface on six squares with permutations $h = (1)(2\,3\,4)(5\,6)$ and $v = (1\,2)(3\,5)(4\,6)$.

In [GJ00] it was shown that when one translation surface is a cover of another, $\pi : (X, \omega) \to (Y, \eta)$, their Veech groups must be commensurable. (In fact, Gabriela Schmithüsen has developed an algorithm for determining the Veech group of a square-tiled surface [Sch04].) Thus (X, ω) is a Veech surface if and only if (Y, η) is a Veech surface. Hence square-tiled surfaces are Veech surfaces since they are covers of tori.

2.3. **Sturmian sequences.** The flat torus $(\mathbb{T}^2, dz)$ can be thought of as a single unit square with its top and bottom edges identified, and its left- and right-hand edges identified as well. If we label the horizontal edge V and the vertical edge H, then a cutting sequence on the torus is a biinfinite sequence of H's and V's.

Remark. We point out that our convention of having V's on the top and bottom of the square and H's on the left- and right-hand sides may seem counter-intuitive, but does offer two small advantages over the more intuitive choice of putting V's on the vertical edges and H's on the horizontal edges.

We will use this same convention to describe labeled edges on square-tiled surfaces, and in that situation crossing an H-edge corresponds to applying the permutation h; and crossing a V-edge corresponds to applying the permutation v, i.e. a side labeled by H is shared between two squares glued by a horizontal translation.

Additionally, with this convention, seeing more H's (respectively, V's) in a cutting sequence means the trajectory is more horizontal (resp. vertical) than vertical (horizontal).

Not all sequences of H's and V's can be obtained as cutting sequences. For example, it is easy to see that the sequence will have isolated V's separated by several H's if the slope of the geodesic is less than 1 as in Figure 2. Having isolated V's (or H's) is not enough to be a cutting sequence, however. If the slope m is less than 1, for example, then the number of H's separating V's may be either $\lfloor \frac{1}{m} \rfloor$ or $\lfloor \frac{1}{m} \rfloor + 1$ with the number of V's in each *block* depending on where the geodesic last intersected the horizontal edge labeled V.

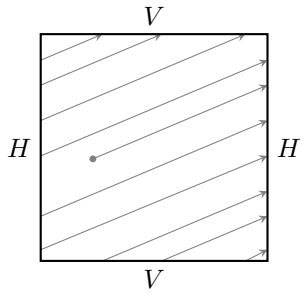

FIGURE 2. The cutting sequence for the geodesic shown here is $HVHHVHHHVH\ldots$.

As tori are Veech surfaces, all geodesics in a given direction will be periodic or uniformly distributed throughout the surface, each case being determined by the slope of the geodesic. If the slope is rational, then the geodesic is periodic. If the slope is irrational, then the geodesic will be uniformly distributed. These two possibilities are reflected in the types of cutting sequences which appear. Periodic geodesics of course give periodic cutting sequences, while uniformly distributed geodesics give *Sturmian sequences*.

A Sturmian sequence is a sequence on two characters that is not eventually periodic and whose complexity function, which counts the number $p(n)$ of distinct subwords of a given length n, has the smallest possible value for a non-eventually periodic sequence: $p(n) = n + 1$. Sturmian sequences are described in detail in [Arn02].

We can determine if a given sequence in the alphabet $E = \{H, V\}$ is a Sturmian sequence, and so a cutting sequence on the torus, by performing a combinatorial operation called *derivation*. As noted above, these sequences must have one of the letters isolated; say V's are isolated for the sake of example. From each block of

consecutive H's we then delete one H obtaining a new sequence which is referred to as the *derived sequence* of the original sequence. If a given sequence is Sturmian, then its derivation will also be Sturmian and so again has an isolated letter. Repeating the process finitely-many times, the other letter H becomes isolated. We can then start the process over again, deleting one letter at a time from blocks of consecutive V's.

Given a Sturmian sequence $\varepsilon \in E^{\mathbb{Z}}$, we can associate a real number by the following procedure. Suppose that the number of H's appearing between adjacent V's is a_0 or a_0+1 (a_0 could be zero). Delete a_0 H's from each block of consecutive H's between a pair of V's. The H's are now the isolated characters. Now suppose between two adjacent H's there are a_1 or a_1+1 V's. Again delete blocks of V's of size a_1, leaving one V if there were a_1+1 V's in the block. Repeat the process alternating between deleting H's and V's where at each step there are a_i or a_i+1 H's (respectively, V's) between the isolated V's (resp., H's). Now define the *slope* of the Sturmian sequence ε to be

$$m = m(\varepsilon) = a_0 + \cfrac{1}{a_1 + \cfrac{1}{a_2 + \cfrac{1}{a_3 + \dots}}}$$

which is slope of the geodesic on the torus with cutting sequence ε.

We will also define the *length* $M(\varepsilon)$ of a Sturmian sequence ε as

$$M = M(\varepsilon) = \left\lfloor \frac{1}{m(\varepsilon)} \right\rfloor,$$

and the *rotation parameter* $\theta(\varepsilon)$ of ε as

$$\theta = \theta(\varepsilon) = \frac{1}{m(\varepsilon)} - M(\varepsilon).$$

If we consider a circle which is a horizontal line on the torus, the first return map of the geodesic flow with slope m is precisely a rotation by θ.

2.4. **Words and languages.** By an *alphabet* $\mathcal{A}$ we mean some finite set of labels. A *word* with letters in $\mathcal{A}$ is a finite, ordered list of elements in $\mathcal{A}$, and the collection of all words with letters in $\mathcal{A}$ is denoted $\mathcal{A}^*$. If we define a binary operation on $\mathcal{A}^*$ by concatenating two words, then $\mathcal{A}^*$ is the free monoid generated by $\mathcal{A}$.

We will use subscripts to denote individual letters of a word $w \in \mathcal{A}^*$: w_0 is the first character of the word, w_1 the second character, and so on. We will denote the length of a word by $|w|$.

In general, a *language* with letters in $\mathcal{A}$ is simply a subset of $\mathcal{A}^*$. We say that a language is *minimal* if for each word w in the language there exists an integer N such that w is a subword of every word in the language of length at least N. Associated to any language is a dynamical system consisting of all sequences where every finite subword of a sequence is a word of the language, together with the shift map. A *cylinder set* $[w]$ is the set of all sequences whose first $|w|$ characters give the word w. These cylinder sets form a base for a topology on the sequence space making it homeomorphic to a Cantor set and for which the shift map is continuous. It is a basic fact that the language is minimal if and only if the associated topological dynamical system is minimal: that is, every orbit is dense in the space. Note that the cylinder sets form a basis for the topology of this space of sequences, and by

the Kolmogorov extension theorem, a Borel measure on this space is determined by the measure of each cylinder.

Given a biinfinite sequence of letters in $\mathcal{A}$, say $u \in \mathcal{A}^{\mathbb{Z}}$, the *language* of u is the collection of all finite subwords which appear in u. We will say the sequence u is *minimal* if its language is minimal.

2.5. **Interval exchange transformations.** An *interval exchange transformation*, abbreviated *IET*, is a right-continuous bijection from an interval I to itself which is a piecewise translation having finitely-many discontinuities. That is, an IET cuts the interval I up into finitely-many pieces and then rearranges the pieces. Such a map is determined by two pieces of information: the *combinatorial data* of the map is the permutation describing the rearrangement of the subintervals, and the *length data* is a vector describing the lengths of the subintervals.

We will represent the combinatorial data of a given interval exchange $T : I \to I$ as a pair of bijections. If $\mathcal{A}$ is a set of k labels, we consider a pair of bijections, π_0 and π_1, from $\mathcal{A}$ to $\{1, 2, ..., k\}$ so that $\pi_0^{-1}(1)$ is the label of the left-most subinterval of I before applying T; $\pi_1^{-1}(1)$ is the label of the left-most subinterval of I after applying T; $\pi_0^{-1}(2)$ is the label of the second interval before applying T, and $\pi_1^{-1}(2)$ is the label of the second interval after applying T; and so on.

We will use the combinatorial data of an interval exchange to define two total orderings on the symbols of $\mathcal{A}$. For each bijection $\pi : \mathcal{A} \to \{1, 2, ..., d\}$ we define $\leq_\pi$ as follows: for $a, b \in \mathcal{A}$, say $a \leq_\pi b$ if $\pi(a) \leq \pi(b)$. By a π-*interval* of $\mathcal{A}$ we mean a collection of consecutive elements of $\mathcal{A}$ according to $\leq_\pi$.

The length data of an interval exchange is given by a vector of lengths $\ell = (\ell_a)_{a \in \mathcal{A}}$ such that each $\ell_a > 0$ and $\sum_a \ell_a$ is the length of the interval I. See Figure 3.

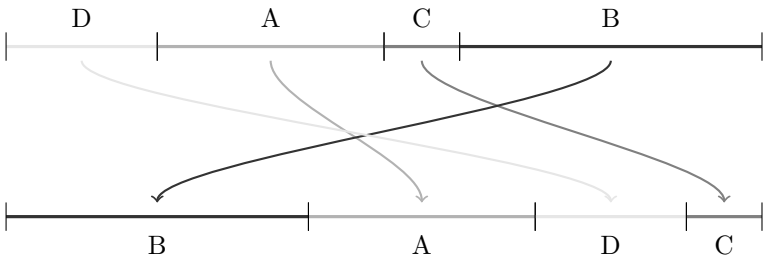

FIGURE 3. The interval exchange with combinatorial data $\pi_0 = (A, B, C, D) \mapsto (2, 4, 3, 1)$, $\pi_1 = (A, B, C, D) \mapsto (2, 1, 4, 3)$ and length data $\ell = (0.2, 0.3, 0.1, 0.4)$.

We say that the combinatorial data (π_0, π_1) is *irreducible* if there does not exist a $1 \leq j < k$ such that

$$\pi_0^{-1}(\{1, 2, ..., j\}) = \pi_1^{-1}(\{1, 2, ..., j\}).$$

If the combinatorial data of an interval exchange was not irreducible, then the interval exchange would decompose into two invariant subintervals for every choice of length data.

Once the combinatorial and length data are known, the actual operation of the map is easy to describe. For each $a \in \mathcal{A}$, let

$$\delta_a^0 = \sum_{b:\pi_0(b)<\pi_0(a)} \ell_b$$
$$\delta_a^1 = \sum_{b:\pi_1(b)<\pi_1(a)} \ell_b$$

and define the subintervals

$$I_a^0 = \left[\delta_a^0, \delta^0_{\pi_0^{-1}(\pi_0(a)+1)}\right)$$
$$I_a^1 = \left[\delta_a^1, \delta^1_{\pi_1^{-1}(\pi_1(a)+1)}\right).$$

Then the δ_a^0 represent the discontinuities of T, and δ_a^1 are the discontinuities of T^{-1}. The map T acts by translating I_a^0 to I_a^1:

$$T(x) = x - \delta_a^0 + \delta_a^1 \text{ if } x \in I_a^0, \text{ and}$$
$$T^{-1}(x) = x - \delta_a^1 + \delta_a^0 \text{ if } x \in I_a^1.$$

We say that an interval exchange T satisfies the *infinite distinct orbit condition*, often abbreviated *idoc*, if for each $n > 0$ and each distinct $a, b \in \mathcal{A}$ with $\pi_0(b) > 1$ we have

$$T^n\left(\delta_a^0\right) \neq \delta_b^0.$$

It was shown in [Kea75] that the interval exchange T is minimal if and only if it satisfies the infinite distinct orbit condition.

Given a finite word $w = w_0 w_1 \cdots w_n \in \mathcal{A}^*$, we will let I_w^0 denote the set of points inside of I which start off in $I_{w_0}^0$, then proceed to $I_{w_1}^0$ when T is applied, then $I_{w_2}^0$ when T^2 is applied, and so on:

$$I_w^0 = \bigcap_{i=0}^{n} T^{-k}\left(I_{w_i}^0\right).$$

Given an interval exchange transformation T and a point $x \in I$, we can consider the biinfinite sequence of labels obtained by recording the label a when an iterate $T^n(x)$ is in the interval I_a^0.

If we are given a sequence $u \in \mathcal{A}^{\mathbb{Z}}$, it is natural to ask if u represents the symbolic trajectory of some point under an interval exchange with intervals labeled by $\mathcal{A}$. Ferenczi and Zamboni give necessary and sufficient conditions in [FZ08] for a sequence of symbols to be the symbolic trajectory of some interval exchange transformation with fixed combinatorics.

Theorem 1 ([FZ08])**.** *The sequence $u \in \mathcal{A}^{\mathbb{Z}}$ is a symbolic trajectory of an interval exchange satisfying Keane's infinite distinct orbit condition with irreducible combinatorial data $\pi_0, \pi_1 : \mathcal{A} \to \{1, 2, ..., k\}$ if and only if the following conditions are satisfied:*

(1) u is a minimal sequence;
(2) each letter of $\mathcal{A}$ appears in u;
(3) for each finite-word w appearing in u, the set of letters which appear as prefixes of w in u forms a π_1-interval;
(4) for each finite-word w appearing in u, the set of letters which appear as suffixes of w in u forms a π_0-interval;

(5) *if $a, b \in \mathcal{A}$ are admissible prefixes of w with $a \leq_{\pi_1} b$, and if y is an admissible suffix of aw and z an admissible suffix of bw, then $y \leq_{\pi_0} z$; and*
(6) *if a and b are admissible prefixes of w, then there is exactly one element of $\mathcal{A}$ which is an admissible suffix of both aw and bw.*

This theorem only says that a given sequence $u \in \mathcal{A}^{\mathbb{Z}}$ is a symbolic trajectory of *some* interval exchange with the chosen combinatorial data, but does not tell us if u is a symbolic trajectory for a particular interval exchange with this combinatorial data. Obtaining a particular interval exchange with the given sequence as a symbolic trajectory is equivalent to finding a shift-invariant measure for the dynamical system associated with the sequence. Given such a measure we can define the length of the subinterval $I_{a_1 a_2 \cdots a_n}$ as the measure of the cylinder $[a_1 a_2 \cdots a_n]$.

3. Interval Exchanges on Square-Tiled Surfaces

Let $\Lambda = \{1, 2, ..., d\}$ be a set of labels, and let h and v be two permutations of Λ such that $\langle h, v \rangle$ acts transitively on Λ. Let X be the corresponding square-tiled surface. Denote the geodesic interval at the base of square λ by I_λ as in Figure 4(a), and let $I = \bigcup_\lambda I_\lambda$. Now consider the first-return map $T_m : I \to I$ obtained by firing a geodesic ray from $x \in I$ with slope m, and recording where the geodesic first intersects I as in Figure 4(b). In Figure 4(b), the base of each I_λ is broken into two pieces: a left-hand side and a right-hand side. Geodesics in a left-hand side of the base of a square stay parallel to one another and are followed until reaching the base of another square. The paths which geodesics in a left-hand interval follow are shaded with various textures in Figure 4(b). Similarly, geodesics emitted from the right-hand side of each base remain parallel and are followed until reaching the base of another square, and the path followed by these right-hand trajectories are given different shades of grey.

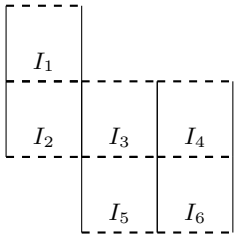

(a) The bases of squares form a collection of disjoint intervals.

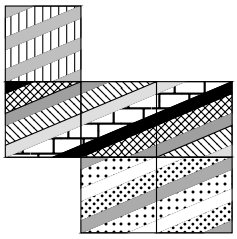

(b) We consider the first-return map on this collection of intervals.

Figure 4. We consider the interval exchanges which are given by the first-return map on the bases of the squares.

In the following we assume that the slope of a geodesic is $m > 0$ and the geodesic flow is to the North-East on the surface.

Lemma 2. *The first-return map on the base of the squares in a square-tiled surface has the form*

$$T_m(x, \lambda) = \begin{cases} (R_\theta(x), vh^M(\lambda)) & \text{if } x \in [0, 1-\theta) \\ (R_\theta(x), vh^{M+1}(\lambda)) & \text{if } x \in [1-\theta, 1) \end{cases}$$

where $\theta \in [0,1)$ and $M \in \{0,1,2,...,D-1\}$ are values depending on m, D depends on the surface, and R_θ is the rotation of $[0,1)$ by θ.

Proof. Since each I_λ is a unit interval, we can think of I as $[0,1) \times \Lambda$, identifying I_λ with $[0,1) \times \{\lambda\}$. Suppose that the horizontal permutation h consists of cycles h_1, h_2, ..., h_k and denote the length of the cycle h_i by $|h_i|$. Define

$$\begin{aligned} D &:= \operatorname{lcm}(|h_1|, |h_2|, ..., |h_k|), \\ M &:= \left\lfloor \frac{1}{m} \right\rfloor \pmod{D}, \\ \theta &:= \frac{1}{m} - \left\lfloor \frac{1}{m} \right\rfloor, \text{ and} \\ R_\theta(x) &:= x + \theta - \lfloor x + \theta \rfloor. \end{aligned}$$

If $x \in I_\lambda$, then the geodesic ray to the North-East emitted from x with slope m cuts across the vertical sides of squares in the same horizontal cylinder as square λ before eventually reaching the base of some square λ' in one of the horizontal cylinders connected to the cylinder containing λ. Because the squares are 1×1 and the geodesic has slope m, the horizontal distance traveled by the geodesic before reaching $I_{\lambda'}$ is $\frac{1}{m}$.

To make the arithmetic simpler, we develop the cylinder containing square λ in the plane, choosing coordinates so that the left-most point of I_λ is $(0,0)$. See Figure 5.

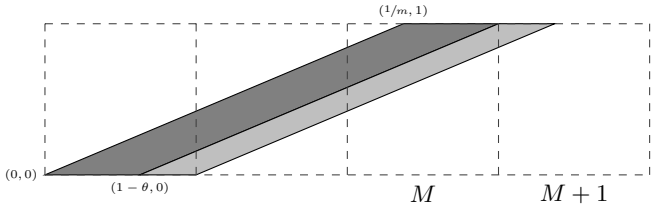

FIGURE 5. Computing the first-return map.

The geodesic emitted from x leaves the cylinder at the point $(y,1)$ where

$$y := x + \frac{1}{m} = x + M + \theta.$$

When $x < 1-\theta$, y is in $[M, M+1)$, and so cuts across M squares before leaving the cylinder and moving up into square $vh^M(\lambda)$. When $x \geq 1-\theta$, $y \in [M+1, M+2)$ and so the geodesic leaves the cylinder by moving into square $vh^{M+1}(\lambda)$.

As we identify the base of square λ' with the interval $[0,1)$, the x-coordinate of our geodesic upon entering I'_λ is

$$x + \frac{1}{m} - \left\lfloor x + \frac{1}{m} \right\rfloor = R_\theta(x).$$

The cycles of h determine horizontal cylinders on the surface where the length of a cylinder equals the length of the corresponding cycle. Given a cycle of length k, then for each square λ in this cylinder, $h^M(\lambda) = h^{M+nk}(\lambda)$ for each $n \in \mathbb{Z}$. Hence we only care about the value of M modulo D. □

Lemma 3. *The first-return map T_m described in Lemma 2 is conjugate to an interval exchange on $[0, d)$.*

Proof. Keeping M and θ as defined in Lemma 2, let T be the interval exchange on $2d$ subintervals of $[0, d)$ with labels from $\mathcal{A} = \Lambda \times \{L, R\}$, combinatorial data

$$\pi_0(\lambda, s) = \begin{cases} 2\lambda - 1 & \text{if } s = L \\ 2\lambda & \text{if } s = R \end{cases}$$

$$\pi_1(\lambda, s) = \begin{cases} 2vh^M(\lambda) & \text{if } s = L \\ 2vh^{M+1}(\lambda) - 1 & \text{if } s = R \end{cases}$$

and length data

$$\ell_{\lambda,s} = \begin{cases} 1 - \theta & \text{if } s = L \\ \theta & \text{if } s = R \end{cases}.$$

With this choice of combinatorial and length data, the discontinuities of T and T^{-1} are respectively

$$\delta^0_{(\lambda,s)} = \begin{cases} \lambda - 1 & \text{if } s = L \\ \lambda - \theta & \text{if } s = R \end{cases} \quad \text{and} \quad \delta^1_{(\lambda,s)} = \begin{cases} vh^M(\lambda) - 1 + \theta & \text{if } s = L \\ vh^{M+1}(\lambda) - 1 & \text{if } s = R. \end{cases}$$

Now consider the map $\varphi : [0, 1) \times \Lambda \to [0, d)$ defined by

$$\varphi(x, \lambda) = x + \lambda - 1$$

It is easy to verify that $\varphi T_m = T\varphi$. If $x \in [0, 1 - \theta)$, then

$$\begin{aligned} \varphi T_m(x, \lambda) &= \varphi(R_\theta(x), vh^M(\lambda)) \\ &= R_\theta(x) + vh^M(\lambda) - 1 \\ &= x + \theta + vh^M(\lambda) - 1 \\ &= x + \lambda - 1 - (\lambda - 1) + vh^M(\lambda) - 1 + \theta \\ &= x + \lambda - 1 - \delta^0_{(\lambda,L)} + \delta^1_{(\lambda,L)} \\ &= T(x + \lambda - 1) \\ &= T\varphi(x, \lambda). \end{aligned}$$

The proof is similar if $x \in [1 - \theta, 1)$. □

In the above we assumed the slope m was positive and the flow was to the North-East just to simplify the statements and proofs of Lemma 2 and Lemma 3. If we considered flowing in another direction, then the above arguments still hold except h and/or v may be replaced by their inverses, and R_θ may be replaced by $R_{-\theta}$. For example, flowing to the North-West with slope $m < 0$ on the surface given by (h, v) is equivalent to flowing to the North-East with slope $-m$ on the surface determined by (h^{-1}, v), and similarly for the other directions.

4. Consistent Sequences and Symmetric Sturmian Sequences

4.1. **Consistent sequences.** Let λ be the label of some square on a square-tiled surface, and consider the geodesic intervals corresponding to the left-hand and bottom edges of λ, which we will label (λ, H) and (λ, V), respectively, as in Figure 6. Now consider the geodesic flow with slope m on the surface, where for simplicity we again consider the case that $m > 0$ and the flow is to the North-East. When the flow enters square λ it must cross either (λ, H) or (λ, V), and then must leave the square by crossing either $(v(\lambda), V)$ or $(h(\lambda), H)$. The order in which the geodesic crosses these edges gives a sequence in $(\Lambda \times E)^{\mathbb{Z}}$. We will adopt the convention that if the geodesic hits the corner of a square which is not a conical singularity, then two symbols are recorded: either $(h(\lambda), H)\,(vh(\lambda), V)$ or $(v(\lambda), V)\,(hv(\lambda), H)$.

If the upper right-hand corner of square λ is not a conical singularity, then $vh(\lambda) = hv(\lambda)$. The squares λ satisfying this property are called *good squares*, and squares with a singularity in the upper right-hand corner are *bad squares*.

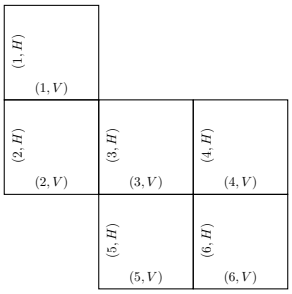

Figure 6. The curves we consider in the cutting sequence.

Our basic question is which sequences in $(\Lambda \times E)^{\mathbb{Z}}$ correspond to cutting sequences on the surface. It is obvious that not all sequences of symbols will be cutting sequences. The simplest property all cutting sequences must satisfy is what we call *consistency*. We say that a sequence $(\lambda_n, \varepsilon_n) \in (\Lambda \times E)^{\mathbb{Z}}$ is *consistent* if $\lambda_{n+1} = h(\lambda_n)$ when $\varepsilon_{n+1} = H$, and $\lambda_{n+1} = v(\lambda_n)$ when $\varepsilon_{n+1} = V$. To simplify notation we will let E act on Λ by $H \cdot \lambda = h(\lambda)$ and $V \cdot \lambda = v(\lambda)$ with $h, v \in \mathfrak{S}_\Lambda$ fixed permutations.

For example, consider the L-shaped surface built from three squares with permutations $h = (1)(2\,3)$ and $v = (1\,2)(3)$ as in Figure 8. A sequence containing

$$...(1, H)\,(2, V)\,(3, H)\,(2, H)...$$

may be consistent (whether it is in fact consistent or not depends of course on the rest of the sequence, but nothing shown here violates consistency), while a sequence containing

$$...(1, H)\,(2, V)\,(1, H)\,(1, H)...$$

can not be consistent because of the $(1, H)$ which follows $(2, V)$. If the pair (λ, H) follows $(2, V)$ then we *must* have $\lambda = 3$ as $h(2) = 3$.

Since a square-tiled surface X is a cover of the square torus $\mathbb{T}^2$, geodesics on X project to geodesics on $\mathbb{T}^2$. Similarly, cutting sequences on X project to cutting sequences on $\mathbb{T}^2$. In particular, if $(\lambda_n, \varepsilon_n)$ is a cutting sequence on the square-tiled surface X, then ε should be a cutting sequence on $\mathbb{T}^2$. We can try to go in the other

direction, lifting a cutting sequence on $\mathbb{T}^2$ to a cutting sequence on X. In particular, given an initial square λ_0 and a cutting sequence ε on $\mathbb{T}^2$, we can construct a *combinatorial lift* by defining $\lambda_{n+1} = \varepsilon_{n+1} \cdot \lambda_n$. We may choose λ_0 to be any initial square, and so on a square-tiled surface built from d squares there are d different combinatorial lifts of each cutting sequence on the torus. Combinatorial lifts are consistent, but not all consistent sequences are combinatorial lifts. Additionally, every cutting sequence on X is a combinatorial lift of some cutting sequence on $\mathbb{T}^2$, but the converse is not true.

4.2. **A combinatorial lift which is not a cutting sequence.** Consider the cutting sequence associated to the geodesic on the square torus with slope $m = \frac{211}{500} = 0.422$ which passes through the point $(\frac{133}{211}, 0)$, where $(0,0)$ is the lower left-hand corner of the square. This geodesic crosses the vertical edge H twice before intersecting the top right-hand corner of the square, but then continues from the bottom left-hand corner giving the cutting sequence $\varepsilon = ...VHHHV...$ as in Figure 7. Here we made the choice that hitting the corner should correspond to recording HV. We could have just as easily decided to instead record VH to obtain the sequence $\varepsilon' = ...VHHVH....$ We will assume that ε_0 and ε'_0 are the first characters that appear in the strings above; that is, $\varepsilon_0 = V$, $\varepsilon_1 = H$, $\varepsilon_2 = H$, $\varepsilon_3 = H$, $\varepsilon_4 = V$, and $\varepsilon'_0 = V$, $\varepsilon'_1 = H$, $\varepsilon'_2 = H$, $\varepsilon'_3 = V$, $\varepsilon'_4 = H$, and so on.

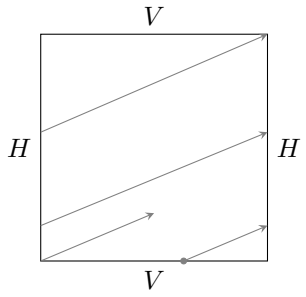

FIGURE 7. On the torus, geodesics are still defined even when they hit the corner.

We will now build a consistent sequence which is not the cutting sequence of a geodesic on a square-tiled surface. We consider the L-shaped surface built from three squares which has permutations $h = (1)(2\,3)$ and $v = (1\,2)(3)$ shown in Figure 8. Now consider the sequences $(\lambda_n, \varepsilon_n)$ and $(\lambda'_n, \varepsilon'_n)$ on $(\{1,2,3\} \times \{H,V\})^{\mathbb{Z}}$ which are defined by $\lambda_0 = \lambda'_0 = 2$ and $\lambda_{n+1} = \varepsilon_{n+1} \cdot \lambda_n$. The sequence λ' is constructed similarly, using ε' in place of ε. We then have

$$(\lambda_n, \varepsilon_n) = ...(2,V)(3,H)(2,H)(3,H)(3,V)...$$
$$(\lambda'_n, \varepsilon'_n) = ...(2,V)(3,H)(2,H)(1,V)(1,H)...$$

By construction both of these sequences are consistent with the permutations defining the surface and are combinatorial lifts of the Sturmian sequences ε and ε', but neither sequence is a cutting sequence on this surface. If $(\lambda_n, \varepsilon_n)$ was the cutting sequence of γ and $(\lambda'_n, \varepsilon'_n)$ the cutting sequence of γ', then γ and γ' would project to the geodesic in Figure 7. However, no infinitely-long geodesic on this L-shaped

surface can project to this geodesic because the geodesic on the torus passes through a corner which is one of the branch points of the cover: on the L-shaped surface this is a conical singularity and the geodesic flow ends upon hitting the singularity.

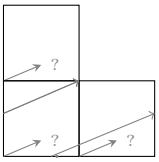

FIGURE 8. We can easily produce sequences of labels which are consistent with the gluings of a square-tiled surface, but which do not represent cutting sequences.

The choice of $\lambda_0 = 2$ is arbitrary in this example. The same phenomenon of constructing a cutting sequence which describes a geodesic which intersects a cone point would happen for any choice of λ_0 on the L-shaped surface since the corner of each square is a cone point. On other surfaces where some corners are cone points and some are not, however, the choice of initial square λ_0 matters. We will see later that cutting sequences which describe geodesics passing through a corner must satisfy a certain symmetry condition, such a sequence is called *almost symmetric* below. These almost symmetric sequences are candidates for cutting sequences of the torus which do not lift to the square-tiled surface, but some may lift and some may not, depending on the choice of λ_0. We will encode the information of which sequences lift and which do not by considering walks on a special type of graph and will see that there is a way of determining precisely which almost symmetric sequences describe geodesic which intersect a cone point according to properties of this walk.

Every geodesic on the square torus $\mathbb{T}^2$ which does not pass through the corner point has d, the number of squares, lifts to a square-tiled surface by basic covering theory. Thus every cutting sequence of a geodesic on $\mathbb{T}^2$, which does not describe a geodesic passing through the corner, has d combinatorial lifts to the square-tiled surface. Hence the only serious obstacle to describing cutting sequences on a square-tiled surface is classifying those sequences which describe geodesics passing through the corner point.

4.3. **Symmetric Sturmian sequences.** We will call a Sturmian sequence ε *odd symmetric* if there exists an N such that $\varepsilon_{N+k} = \varepsilon_{N-k}$, and *even symmetric* if there exists an N such that $\varepsilon_{N+k} = \varepsilon_{N-k-1}$ for every $k \geq 0$. We will say ε is *almost symmetric* if there exists an N such that $\varepsilon_{N+k} = \varepsilon_{N-k-1}$ for every $k \geq 1$ and $\varepsilon_N \neq \varepsilon_{N-1}$. Below is an example of an odd symmetric sequence, an even symmetric sequence, and two almost symmetric sequences. The points around which the sequences are symmetric are highlighted.

$...HVVHVVV\mathbf{H}VVVHVV...$	Odd symmetric
$...VVHVVH\mathbf{VV}HVVHVV...$	Even symmetric
$...HVVHVV\mathbf{VH}VVHVVV...$	Almost symmetric
$...HVVHVV\mathbf{HV}VVHVVV...$	Almost symmetric

By a *symmetric* Sturmian sequence we will mean a sequence which is odd symmetric, even symmetric, or almost symmetric.

Recall that the torus is a hyperelliptic surface and the hyperelliptic involution $\iota : \mathbb{T}^2 \to \mathbb{T}^2$ acts as a 180°-rotation of the square about the center point. There are four Weierstrass points fixed by ι indicated in Figure 9.

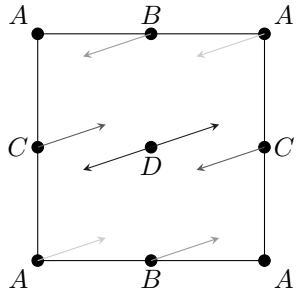

FIGURE 9. The four Weierstrass points on the torus together with geodesics through those points.

Lemma 4. *A Sturmian sequence is symmetric if and only if it represents the cutting sequence of a uniformly distributed geodesic which passes through a Weierstrass point of the torus. Furthermore,*

(i) even symmetric Sturmian sequences correspond to geodesics passing through the center point D;
(ii) almost symmetric sequences pass through the corner point A.
(iii) odd symmetric sequences correspond to geodesics passing through B or C; and

Recall that we are considering uniformly distributed geodesics which must have irrational slope, and such a geodesic can not pass through two distinct Weierstrass points.

Proof of Lemma 4. Let γ denote the geodesic whose cutting sequence is the Sturmian sequence ε. Suppose that γ passes through the point P on the side of the square labeled ε_0. Applying ι thus gives a geodesic $\iota(\gamma)$ which passes through the point $\iota(P)$ which must also be on the side of the square labeled ε_0, since ι preserves the sides of the square. Furthermore, the cutting sequence $\iota(\gamma)$ is the same as the cutting sequence of γ, provided we walk along the geodesics γ and $\iota(\gamma)$ according to how those geodesics are oriented. If we had adopted the convention to always walk along the geodesic in the positive (North-East) direction, or always in the negative (South-West) direction, then the cutting sequences would reverse. When γ passes through a Weierstrass point, however, the geodesic γ is preserved and we see the same cutting sequences, emitted from that Weierstrass point, in both the positive and negative direction, giving the symmetries as described above. The only oddity here is to recall that when the Weierstrass point under consideration is the corner point we record the symbols VH, or HV, and this results in the "almost symmetric" sequence. This establishes one direction of the lemma, and for the converse we consider the three cases of even symmetric, odd symmetric, and almost symmetric separately.

In each case we have the following setup: supposing ε is symmetric, there exist points P and Q on $\gamma \cap (H \cup V)$ such that the geodesic emitted from P in the positive direction has the same cutting sequence as the geodesic emitted from Q in

the negative direction. Equivalently, the geodesic emitted from $\iota(Q)$ in the positive direction gives the same cutting sequence as the geodesic emitted from P in the positive direction.

(i) Suppose ε is even symmetric, so the ray emitted from P in the positive direction is $\varepsilon' = \varepsilon_0\varepsilon_1\varepsilon_2\cdots$, and the ray emitted from Q in the negative direction is $\varepsilon' = \varepsilon_{-1}\varepsilon_{-2}\varepsilon_{-3}\cdots$ with $\varepsilon_{-1} = \varepsilon_0$, $\varepsilon_{-2} = \varepsilon_1$ and so on. As the ray emitted from P gives $\varepsilon_{\geq 0}$ while the ray emitted from Q gives $\varepsilon_{<0}$, the geodesic γ connects Q to P without first intersecting $H \cup V$. Additionally, since $\varepsilon_0 = \varepsilon_{-1}$, P and Q must both be on the same edge, H or V. Since P and Q are involutes, however, the only geodesic that connects them without first passing through the sides H or V must pass through the center point D of the square.

(ii) Suppose ε is almost symmetric, with the ray emitted from P being $\varepsilon_{\geq 1}$ and the ray emitted from Q being $\varepsilon_{<-1}$ with $\varepsilon_i = \varepsilon_{-i-1}$ for $i > 0$, and $\varepsilon_{-1}\varepsilon_0$ being one of HV or VH.

The portion of the geodesic from Q to P is responsible for the HV (or VH) at $\varepsilon_{-1}\varepsilon_0$. Suppose we followed γ backwards from P to some point R on the vertical edge H. We would likewise follow γ forwards from Q to $\iota(R)$ which would also be on H. If $R \neq \iota(R)$, then this would produce two distinct H's, which do not appear at positions ε_{-1} and ε_0 in our sequence. Similarly, if R were on the horizontal edge V, then $\iota(R)$ would be as well, and if $R \neq \iota(R)$ we would have two distinct V's which do not appear at positions ε_{-1} and ε_0 in the sequence. We must then have that $R = \iota(R)$. If R was either of the Weierstrass points labeled B or C in Figure 9 we would thus produce a single character in our sequence, but not the two characters we lack.

Since our sequence contains VH (or HV) and intersecting a non-Weierstrass point would produce HH or VV, and a Weierstrass point which is not A would produce a since H or single V, the only remaining option is for the geodesic to pass through the Weierstrass point labeled A at the corner of the square to produce the VH (or HV).

(iii) If ε is odd symmetric, then we may choose P and Q such that the ray emitted from P gives $\varepsilon_{>0}$ while the ray emitted from Q gives $\varepsilon_{<0}$ with $\varepsilon_i = \varepsilon_{-i}$ for $i > 0$. This again implies that P and Q are both on H or both on V.

Following γ from Q to P crosses the edge ε_0, while following γ backwards from P to Q crosses ε_0 at the same point. Since Q and P are involutes, this point on ε_0 is fixed, and so is either A, B, or C. The point can not be A, however, as geodesics through A give almost symmetric sequences, and our sequence is not almost symmetric. Thus an odd symmetric sequence must pass through either B or C.

□

Recall that given a Sturmian sequence ε, its *derived sequence* is the sequence ε' obtained from ε by deleting one character from each block of the non-isolated character. Geometrically, this derivation corresponds to applying the vertical, downward Dehn twist when V is isolated, or the horizontal, left-ward Dehn twist when H is isolated. This alters the direction of a geodesic, and the derived sequence is simply the cutting sequence of this new geodesic.

Lemma 5. *If a Sturmian sequence ε is almost symmetric, then so is its derivation ε'.*

Proof. If ε is almost symmetric, then it is the cutting sequence of a geodesic through the corner of the square. Since the corners are preserved by the Dehn twists above, the derivation ε' also corresponds to a geodesic through the corner and so is almost symmetric as well. □

5. Characterizing Cutting Sequences

In this section we prove a characterization theorem for cutting sequences on square-tiled surfaces. We will first make note of certain necessary conditions such a cutting sequence must satisfy, and then convert a sequence satisfying those conditions into a symbolic trajectory for an interval exchange and verify the Ferenczi-Zamboni conditions. We then show that there exists a shift-invariant measure agreeing with the lengths of intervals on our square-tiled surface to conclude that the sequence is a symbolic trajectory for the interval exchange on our surface.

5.1. **Necessary conditions.** Recall that a sequence $(\lambda_n, \varepsilon_n)$ in letters $\Lambda \times E$ is *consistent* if $\lambda_{n+1} = \varepsilon_{n+1} \cdot \lambda_n$ and is a *combinatorial lift* if ε_n is the cutting sequence of a geodesic on $\mathbb{T}^2$. Obviously being a combinatorial lift is a necessary condition for $(\lambda_n, \varepsilon_n)$ to be a cutting sequence, but as the example in Section 4.2 shows it is in general not a sufficient condition.

We will say that $(\lambda_n, \varepsilon_n)$ is *almost symmetric around a bad square* if ε_n is an almost symmetric sequence, so $\varepsilon_{N+k} = \varepsilon_{N-k-1}$ for some N and all $k \geq 1$ but $\varepsilon_N \neq \varepsilon_{N-1}$, and λ_N is a bad square.

5.2. **The periodic case.** If $(\lambda_n, \varepsilon_n)$ is a periodic combinatorial lift, then ε_n must be periodic. Periodic cutting sequences on the torus correspond to geodesics of rational slope on the torus. Any lift of such a geodesic to the square-tiled surface with permutations h and v is a periodic geodesic as well by basic lifting properties of covering spaces and the fact that the square-tiled surface and torus are locally isometric.

5.3. **The non-symmetric case.** Combinatorial lifts of Sturmian sequences which are not almost symmetric, and so do not represent cutting sequences passing through a corner of the square, easily lift to cutting sequences on every square-tiled surface.

It is clear, however, that there may be geodesics on a square-tiled surface whose projection to the torus gives an almost symmetric cutting sequence. This can happen if the geodesic goes through the corner of a good square. To determine when this happens we will associate a family of graphs to the surface which describe possible transitions between the squares when the geodesic leaves one horizontal cylinder for another.

5.4. **The Γ^M graphs.** In the description of T_m from Section 3, the left-hand subinterval of I_λ always maps to the right-hand side of some $I_{\lambda'}$ and similarly the right-hand subinterval maps to the left-hand side of some $I_{\lambda''}$ where $\lambda'' = vhv^{-1}(\lambda')$. This is true for every slope m, and so the only part of the combinatorial data of the interval exchange that may change is the number of squares to cross over before a point on the base exits the top of a cylinder. See Figure 5 on page 9. Let M denote the number of vertical edges (labeled H) crossed when transitioning from the left-hand side of I_λ to the right-hand side of $I_{\lambda'}$ In Figure 5, for example, $M = 2$. In terms of Sturmian sequences, M is the minimal number of H's between

two V's in the cutting sequence, regardless of which letter is isolated. The value of M is determined by the slope m, and considering that the line emitted from $(0,0)$ with slope m intersects the line $y=1$ at the point $(1/m, 1)$, it is easy to see that $M = \lfloor \frac{1}{m} \rfloor$.

Notice that if V's are isolated (meaning $m < 1$), then the blocks of H's between V's have length M or $M+1$. If H's are isolated ($m > 1$, and so $M = 0$), then the number of H's between V's is either 0 or 1. As blocks of consecutive H's may have length M or $M+1$, we will refer to M as the *length* of the Sturmian sequence.

The interval exchange T_m is completely determined by the slope m; in particular, writing $\frac{1}{m} = M + \theta$, the value of θ determines the length data of T_m and the value of M determines the combinatorial data. Representing this combinatorial data as a pair of permutations (π_0, π_1) for a given T_m the permutation π_0 is the same for all values of M while π_1 may change. For this reason we will write π_1^M to mean the permutation associated to the interval exchange T_m when $\lfloor \frac{1}{m} \rfloor = M$.

The combinatorial information of (π_0, π_1^M) can be described as a graph $\Gamma^M = \Gamma^M_{h,v}$. This will be a labeled, directed multigraph with vertex set Λ, each vertex having two outgoing edges labeled L and R, and two incoming edges labeled L and R. We will call such a multigraph a *2-oriented graph*. The edges from vertex λ in graph Γ^M are given by

$$vh^M(\lambda) \xleftarrow{\ L\ } \lambda \xrightarrow{\ R\ } vh^{M+1}(\lambda).$$

If λ belongs to a cycle of h with length ℓ, then $h^M(\lambda) = h^{M+k\ell}$ for all k. Letting h_1, h_2, ..., h_p denote the cycles of h, and $|h_i|$ the length of cycle h_i, there are only

$$D = \operatorname{lcm}(|h_1|, |h_2|, ..., |h_p|)$$

possible choices of combinatorial data (π_0, π_1^M). To see this, notice that if $M > D$, say $M = pD + r$ for some $0 \le r < M$, then for each cycle h_i we have

$$h_i^M = h_i^{pD+r} = \left(h_i^D\right)^p \cdot h_i^r = h_i^r$$

since $h_i^D = \mathrm{id}$ as D is a multiple of the length of cycle h_i. This is true for each cycle h_i and so $h^M = h^r$. As T_m takes points on the left-hand side of the base of square λ to square $vh^M(\lambda)$, and points on the right-hand side of the base of square λ to square $vh^{M+1}(\lambda)$, the value of M only matters modulo D, and so there are only D choices for the combinatorial data of T_m.

As an example, consider the flow on the surface of Figure 1 in a direction with slope m in the interval $(1/3, 1/2)$. For such a slope we will have $M = \lfloor \frac{1}{m} \rfloor = 2$ and so consider the graph Γ^2 where for the vertex labeled λ has an outgoing edge to $vh^2(\lambda)$ labeled L, and an outgoing edge to $vh^3(\lambda)$ labeled R as in Figure 10.

Each edge of this graph represents a possible transition in symbolic trajectories of T_m: the edge labeled L out of λ and ending at $vh^M(\lambda)$ corresponds to points in the $[0, 1-\theta)$-portion of I_λ being mapped to $I_{vh^M(\lambda)}$, and similarly for the edge labeled R. Hence every symbolic trajectory of T_m determines an infinite walk on Γ^M. We will relate cutting sequences to symbolic trajectories of the interval exchanges, and so infinite walks on these graphs.

5.5. **Combinatorial lifts and symbolic trajectories.** Let $(\lambda_n, \varepsilon_n)$ be a combinatorial lift of some Sturmian sequence. We will convert $(\lambda_n, \varepsilon_n)$ to a sequence which we will show is the symbolic trajectory of an interval exchange on the square-tiled surface. First we define a subsequence $\mu_i = \lambda_{n_i}$ of (λ_n) by deleting those λ_n

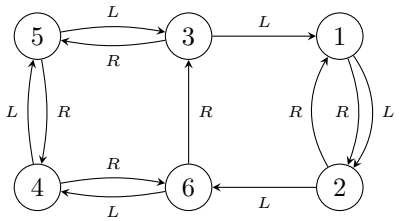

FIGURE 10. The graph Γ^2 associated to the square-tiled surface with permutations $h = (1)(2\,3\,4)(5\,6)$ and $v = (1\,2)(3\,5)(4\,6)$.

where $\varepsilon_n = H$. Now define a sequence $(\sigma_i) \in \Sigma^{\mathbb{Z}} = \{L, R\}^{\mathbb{Z}}$ by considering the distance from λ_{n_i} and $\lambda_{n_{i+1}}$:

$$\sigma_i = \begin{cases} L & \text{if } n_{i+1} - n_i = M + 1 \\ R & \text{if } n_{i+1} - n_i = M + 2 \end{cases}$$

where M is the length (i.e., minimal number of H's between two V's – this may be zero) of the Sturmian sequence ε.

We claim this new sequence $(\mu_n, \sigma_n) \in (\Lambda \times \Sigma)^{\mathbb{Z}}$ is a symbolic trajectory for an interval exchange with combinatorial data (π_0, π_1^M). To see this we need to verify the Ferenczi-Zamboni conditions. We will establish this through a series of lemmas describing properties of (μ_n, σ_n) and the graph Γ^M which encodes the combinatorial data.

Lemma 6. *If $(\lambda_n, \varepsilon_n)$ is a combinatorial lift of a Sturmian sequence ε, then the sequence σ defined above is a Sturmian sequence in the letters L and R. If ε is almost symmetric, then so is σ.*

Proof. Let γ be the geodesic on the square torus represented by the Sturmian sequence ε, let I be the interval at the base of this square, and $R : I \to I$ the circle rotation (interval exchange transformation) given by the first-return map to the base I of the square under the geodesic flow in the direction of γ. This is an interval exchange on two intervals; we label the first, left-hand interval L, and the second, right-hand interval R. The sequence σ is precisely the sequences of L's and R's obtained by walking along γ and recording an L when γ intersects the left-hand interval and an R when γ intersects the right-hand interval. Since this sequence of L's and R's records the orbit of a point under an irrational rotation (irrational since ε is Sturmian hence the slope of the geodesic is not a rational number), σ is itself a Sturmian sequence. See Figure 11 which shows the two different, parallel geodesic segments. (The picture is drawn on a portion of the universal cover of the torus to make the picture easier to understand.) Notice that the geodesic γ_L which passes through the left-hand interval at the base of the first square intersects two vertical lines, while the geodesic γ_R which passes through the right-hand interval crosses three vertical lines. That is, the corresponding cutting sequence is $VHHV$ for γ_L and $VHHHV$ for γ_R. In constructing σ we are simply using the number of H's (vertical line segments crossed) to determine if the geodesic previously passed through a left- or right-hand interval on the base of the square.

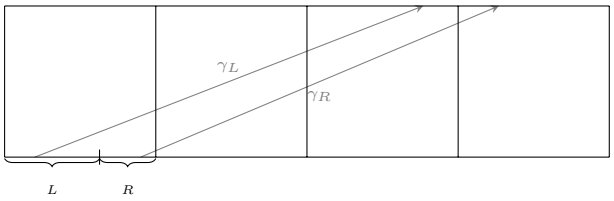

FIGURE 11. The distance between two consecutive V's in the cutting sequence of a geodesic determines, and is determined by, whether the geodesic enters the left- or right-hand side of the base of the square.

If V is the isolated character in an almost symmetric Sturmian sequence ε, then the sequence has the form

$$\overline{\omega}HV^MVHV^MH\omega \quad \text{or} \quad \overline{\omega}HV^MHVV^MH\omega.$$

with $\overline{\omega}$ being the same sequence as ω, but in reverse order. The constructed sequence σ thus has the form

$$\overline{\tau}RL^MRLL^MR\tau \quad \text{or} \quad \overline{\tau}RL^MLRL^MR\tau.$$

where τ is obtained by applying the rule for determining σ_i above to the sequence ω.

Similarly, if ε is an almost symmetric Sturmian sequence but with isolated character V, then ε has the form

$$\overline{\omega}VH^MVHH^MV\omega \quad \text{or} \quad \overline{\omega}VH^MHVH^MV\omega,$$

and the corresponding σ sequence has the form

$$\overline{\tau}LR\tau \quad \text{or} \quad \overline{\tau}RL\tau.$$

In either case, we exactly have that σ is almost symmetric. □

Lemma 7. *Each Γ^M graph is strongly connected.*

Proof. As the edges from vertex λ go to $vh^M(\lambda)$ and $vh^{M+1}(\lambda)$, we must show that there is some composition of powers of vh^M and vh^{M+1} which sends any given λ_i to any λ_j. Since $\langle h, v\rangle$ acts transitively on Λ, it suffices to show that h and v can be formed by compositions of powers of vh^M and vh^{M+1}.

Since Λ has finitely-many elements, there exists some $p > 0$ such that $(vh^M)^p = \mathrm{id}$ and so $(vh^M)^{p-1} = (vh^M)^{-1}$. Thus

$$(vh^M)^{p-1}vh^{M+1} = h^{-M}v^{-1}vh^{M+1} = h.$$

Similarly, there exists a q such that $h^q = h^{-1}$ and so

$$vh^M((vh^M)^{p-1}vh^{M+1})^{qM} = vh^M(h^q)^M = vh^Mh^{-M} = v.$$

□

Corollary 8. *The permutations (π_0, π_1^M) are irreducible.*

Proof. If the permutations were not irreducible, then the corresponding interval exchange on the surface would have a proper subset of intervals which are left invariant by the interval exchange. These invariant subsets of intervals would correspond to connected components of the graph Γ^M, but Lemma 7 shows there is only one such component. □

We will say that an edge of Γ^M is a *bad edge* if traversing that edge could correspond to hitting a conical singularity on the square-tiled surface. Each bad square of the surface gives two bad edges on each Γ^M: if λ is a bad square, then the edge labeled L with target $v(\lambda)$ and the edge labeled R with target $vh(\lambda)$ are bad edges.

Notice that $(\mu_n, \sigma_n)_{n\in\mathbb{Z}}$ determines a walk on the graph Γ^M, where μ_{n+1} is obtained by traveling from μ_n along edge σ_n. We will say that the walk (μ_n, σ_n) is *almost symmetric about a bad edge* if there exists an N such that $\sigma_{N+k} = \sigma_{N-k-1}$ for each $k \geq 1$ and the corresponding edge

$$\mu_N \xrightarrow{\sigma_N} \mu_{N+1}$$

is a bad edge of Γ^M.

We now claim that *Sturmian walks*, walks where the sequence of crossed edges form a Sturmian sequence, on a Γ^M graph are in one-to-one correspondence with symbolic trajectories of interval exchanges T_m on the corresponding square-tiled surface.

Notice that symbolic trajectories of the interval exchange on the bases of squares in a square-tiled surface give Sturmian walks on the corresponding Γ^M graph. This is simply because the sequences of edges labeled L and R coincide with the moments when the corresponding geodesic on the torus intersects the left- and right-hand intervals on the base of the square. As we are considering irrational directions, these L and R subintervals are given by an irrational rotation, and so produce a Sturmian sequence. To show that each such Sturmian walk corresponds to a symbolic trajectory of the interval exchange we need to verify the Ferenczi-Zamboni conditions. We first introduce some notation for the admissible prefixes and suffixes of a word.

Given a word $w = (\nu_0, s_0)(\nu_1, s_1)\cdots(\nu_k, s_k) \in (\Lambda \times \Sigma)^*$, let $(p_n(w))_{n\in\mathbb{Z}}$ be the sequence which gives the locations of w in the sequence $(\mu_n, \sigma_n) \in (\Lambda \times \Sigma)^{\mathbb{Z}}$. That is, $p_n(w)$ satisfies

$$\begin{aligned}(\mu_{p_n(w)}, \sigma_{p_n(w)}) &= (\nu_0, s_0),\\ (\mu_{p_n(w)+1}, \sigma_{p_n(w)}) &= (\nu_1, s_1),\\ &\vdots\\ (\mu_{p_n(w)+k}, \sigma_{p_n(w)}) &= (\nu_k, s_k).\end{aligned}$$

For example, suppose $\Lambda = \{1, 2\}$ and consider the word

$$w = (1, L)(2, R)(1, L)(2, R).$$

If the sequence (μ_n, σ_n) contained

$$...(2, L)(1, L)(2, R)(1, L)(2, R)(1, L)(2, R)(1, L)(2, R)(1, R)...$$

with the initial $(2, L)$ corresponding to (μ_0, σ_0), then we would have $p_0(w) = 1$, $p_1(w) = 3$, $p_2(w) = 5$. Notice that we allow different copies of w to overlap one another in our definition of $p_n(w)$.

We will adopt the convention that $p_0(w)$ is the first occurrence of w starting at an index $n \geq 0$ in (μ_n, σ_n), though this particular choice of starting point will not affect what is to follow.

Following the notation of [FZ08], we let $A(w)$ denote the set of characters which occur immediately before w appears in (μ, σ), and $D(w)$ denotes the set of characters which occur immediately after w:

$$A(w) = \left\{(\ell, s) \,\middle|\, \ell = \mu_{p_n(w)-1},\ s = \sigma_{p_n(w)-1} \text{ for some } n\right\} \text{ and}$$
$$D(w) = \left\{(\ell, s) \,\middle|\, \ell = \mu_{p_n(w)+k+1},\ s = \sigma_{p_n(w)+k+1} \text{ for some } n\right\}.$$

In the example above we would thus have $(2, L), (2, R) \in A(w)$ and $(1, L), (1, R) \in D(w)$. In principle the set of prefixes and suffixes, which were called *arrival* and *departure* sets in [FZ08], for a general language could be quite large and complicated sets of letters. The following lemma shows that these sets are actually very simple for the languages we are considering.

Lemma 9. *The sets $A(w)$ and $D(w)$ each contain at most two characters, $A(w)$ is a π_1^M-interval, and $D(w)$ is a π_0-interval.*

Proof. It is clear that $A(w)$ has two characters since there are only two edges in Γ^M ending at ν_0: one labeled L and one labeled R. From the definition of the graph, we can see that the edge labeled L ending at ν_0 starts at vertex $h^{-M}v^{-1}(\nu_0)$ and the edge labeled R ending at ν_0 starts at vertex $h^{-(M+1)}v^{-1}(\nu_0)$. It is also easy to see from the definition of the permutation π_1^M that these two symbols, $(h^{-M}v^{-1}(\nu_0), L)$ and $(h^{-(M+1)}v^{-1}(\nu_0), R)$, are π_1^M-consecutive:

$$\pi_1^M(h^{-M}v^{-1}(\nu_0), L) = 2vh^M h^{-M}v^{-1}(\nu_0) = 2\nu_0$$
$$\pi_1^M(h^{-(M+1)}v^{-1}(\nu_0), R) = 2vh^{M+1}h^{-(M+1)}v^{-1}(\nu_0) = 2\nu_0 - 1.$$

The proof that $D(w)$ is a π_0-interval and that it contains at most two elements is similar. □

Lemma 10. *If (σ_n) is a Sturmian sequence in $\{L, R\}$, any (μ_n, σ_n) satisfying $\mu_{n+1} = \sigma_n \cdot \mu_n$ is minimal.*

Proof. We need to show that for each finite word of (μ_n, σ_n), there exists some bound M such that every word of length M has the given finite word as a factor. Shifting if necessary, we may assume the finite word is

$$(\mu_0, \sigma_0)\,(\mu_1, \sigma_1)\,\cdots(\mu_k, \sigma_k).$$

As (σ_n) is a Sturmian sequence, there exists some N_0 such that all words of (σ_n) of length N_0 have $\sigma_0\,\sigma_1\,\cdots\,\sigma_k$ as a factor. Letting (τ_i) be the sequence of integers satisfying

$$\sigma_{\tau_i} = \sigma_0,\ \sigma_{1+\tau_i} = \sigma_1,\ \cdots,\ \sigma_{k+\tau_i} = \sigma_k,$$

with $\tau_0 = 0$, the (τ_i) sequence has bounded gaps: $\tau_{i+1} - \tau_i < N_0$. To show minimality of (μ_n, σ_n), it suffices to show that there exists an M such that $\mu_{\tau_i} = \mu_0$ for some $i < M$. To show this it will be helpful to think about the sequence (μ_n, σ_n) as determining a walk on a graph Γ^M.

Let ψ_0 be the empty word and set $\psi_i = \sigma_{\tau_i - 1}\, \sigma_{\tau_i - 2} \cdots \sigma_{\tau_{i-1}}$. We can interpret the ψ_i's both as words of (σ_n), and as permutations of Λ with ψ_0 being the identity. Notice $\mu_{\tau_1} = \psi_1(\mu_0)$ and $\mu_{\tau_i} = \psi_i(\mu_{\tau_{i-1}}) = \psi_i\, \psi_{i-1} \cdots \psi_1(\mu_0)$.

The first few edges crossed by the walk (μ, σ) are as follows:

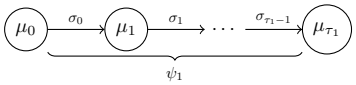

where the walk from μ_0 to μ_{τ_1} has at most N_0 characters. If $\mu_{\tau_1} \neq \mu_0$, then consider a longer portion of the walk until ψ_1 appears again, ending at some $\mu_{\tau_{i_1}}$. Since (σ_n) is Sturmian, there exists some N_1 so that all words of length N_1 have ψ_1 as a factor. Thus the distance from μ_0 to $\mu_{\tau_{i_1}}$ is at most $N_0 + N_1$. If μ_0 does not appear in μ_{τ_1}, μ_{τ_2}, ..., $\mu_{\tau_{i_1}}$, notice that $\mu_{\tau_{i_1}}$ can not equal μ_0 (by assumption), but can also not equal μ_{τ_1} since $\mu_{\tau_{i_1}} = \psi_1(\mu_{\tau_{(i_1 - 1)}})$ and $\mu_{\tau_{(i_1-1)}} \neq \mu_0$.

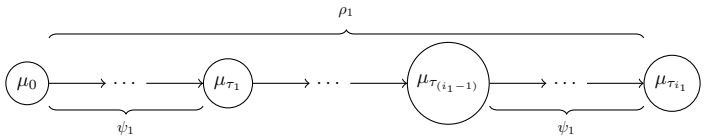

Let ρ_1 be the product of all the symbols (permutations) between μ_0 and $\mu_{\tau_{i_1}}$. As (σ_n) is Sturmian, there exists an N_2 such that all words of length N_2 contain ρ_1 as a factor. Consider now the shortest portion of our walk which contains a second factor of ρ_1 and ends at some $\mu_{\tau_{i_2}}$.

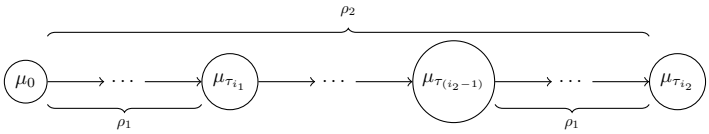

If no μ_1 through $\mu_{\tau_{i_2}}$ equals μ_0, then $\mu_{\tau_{i_2}}$ is not equal to μ_0, μ_{τ_1}, or $\mu_{\tau_{i_1}}$, and the number of steps from μ_0 to $\mu_{\tau_{i_2}}$ is at most $N_0 + N_1 + N_2$. Letting ρ_2 denote the transitions from μ_0 to $\mu_{\tau_{i_2}}$, there exists some N_3 such that all words of length at least N_3 contain a factor of ρ_2, and we can consider the shortest portion of our walk starting from μ_0 and containing two disjoint factors of ρ_2.

We continue in this way building a sequence of symbols $\mu_{\tau_{i_k}}$ which can not take on $k+1$ different symbols of Λ. If we repeat this procedure up until $\mu_{\tau_{i_d}}$, then this symbol can not take on any value in the alphabet. Hence at some point μ_0 must repeat, and this must happen within $N_0 + N_1 + N_2 + \cdots + N_d$ steps. □

Corollary 11. *If $(\lambda_n, \varepsilon_n)$ is a combinatorial lift of a Sturmian sequence, then the associated sequence (μ_n, σ_n) is minimal.*

Lemma 12. *If (σ_n) is a Sturmian sequence in $\Sigma^{\mathbb{Z}}$, each (μ_n, σ_n) satisfying $\mu_{n+1} = \sigma_n \cdot \mu_n$ contains every element of $\Lambda \times \Sigma$.*

Proof. Thinking of (μ_n, σ_n) as a walk on a graph $\Gamma = \Gamma^M$, the lemma claims that every edge of the graph is crossed by the walk. Suppose this was not the case and there was some edge the walk does not cross. We will perform an iterative procedure replacing our graph with a *derived graph* Γ' so that the uncrossed edge

of Γ becomes a vertex which the corresponding *derived walk* (μ'_n, σ'_n) avoids. The derived graph will be a 2-oriented graph with the same number of vertices and edges as Γ, and so a walk which avoids one vertex implies there are at least three edges of Γ' which the derived walk does not cross. Repeating the derivation procedure produces a graph Γ'' and a walk (μ''_n, σ''_n) which avoids two vertices, and at least four edges are not crossed. Since our graphs all have the same finite number of vertices and edges, this procedure will after finitely-many steps show that our walk must not cross *any* edge.

We define the derived graph Γ' and derived walk (μ'_n, σ'_n) as follows. Suppose that the edge of Γ which is not crossed is

$$\lambda \xrightarrow{\sigma} \lambda'.$$

Let τ denote the element of Σ which is not the σ above. Now consider blocks of (σ_n) which consist of blocks of consecutive τ's and end in σ. There will be two such blocks, but the possible blocks depend on the length M of the Sturmian sequence and whether σ or τ is the isolated character. We will refer to the two possible blocks as L' and R'.

If σ is the isolated character of (σ_n), then the blocks ending in σ have the form

$$L' = \tau^M \sigma \text{ or } R' = \tau^{M+1}\sigma.$$

If τ is the isolated character, then the blocks have the form

$$L' = \sigma \text{ or } R' = \tau\sigma.$$

In either case we may interpret L' and R' as permutations of Λ and consider the graph Γ' with vertices the characters of Λ and edges

$$L'(\lambda) \xleftarrow{L'} \lambda \xrightarrow{R'} R'(\lambda).$$

Notice that Γ' is a 2-oriented graph with the same number of edges and vertices as Γ. Furthermore, the proof of Lemma 7 applies equally well to Γ', and so Γ' is strongly connected.

The derived walk (μ'_n, σ'_n) consists of the sequence σ'_n of L' and R' obtained by replacing the blocks in (σ_n) above with L' and R', and setting $\mu'_0 = \mu_0$ and $\mu'_{n+1} = \sigma'_n \cdot \mu'_n$.

Notice that (σ'_n) is a Sturmian sequence. In the case that τ is the isolated character, replacing blocks of $\tau\sigma$ with R' and blocks of σ with L' is the same as deriving the Sturmian sequence (removing a single σ from blocks of consecutive σ's), which gives a Sturmian sequence, and then doing a character-by-character replacement by substituting τ with R' and σ with L', which is obviously a Sturmian sequence. In the case that σ is the isolated character, (σ'_n) is obtained by deriving (σ_n) M times so that τ becomes the isolated character, and then performing the operation previously described.

The derived walk (μ'_n, σ'_n) on Γ' is an acceleration of the walk (μ_n, σ_n) on Γ where we consider longer pieces of the original walk ending in σ. Since $\lambda \xrightarrow{\sigma} \lambda'$ is not crossed in the original walk, there are no blocks crossing such an edge. Since the other incoming edge around λ' must be labeled by τ, the derived walk avoids the vertex λ' of Γ'. As there are two edges which enter λ', the derived walk must not cross either of these edges, and can not cross either of the outgoing edges either. It may happen that one of the incoming edges may also be an outgoing edge (i.e., there may be a loop at λ'). Thus there are at least three forbidden edges of Γ' the

derived walk can not cross. There can not be fewer than three edges because this would violate the strong connectedness of the graph.

As (σ'_n) is a Sturmian sequence, one character, say σ', is isolated and we can repeat the above procedure by considering blocks of (σ'_n) which end in σ' to produce a second derived graph Γ'' and second derived walk (μ''_n, σ''_n). Again, Γ'' will be a strongly-connected 2-oriented graph and (σ''_n) will be a Sturmian sequence. As (σ'_n) avoids two edges, the walk (μ''_n, σ''_n) avoids at least four edges. Loops may occur and so the number of avoided edges does not double each time. However, the strong connectedness of the graph prevents us from having a connected component of avoided edges and so the number of forbidden edges must grow by at least one.

Each time we replace a graph and a walk on that graph with its derivation in this manner, the number of avoided edges grows by at least one. Since there are only finitely-many edges, however, after finitely-many derivations, our walk, which is given by a Sturmian sequence, can not cross any edge of the graph. Thus there can not be any edge avoided by the graph, and so the original sequence (μ_n, σ_n) must contain each element of $\Lambda \times \Sigma$. □

Proposition 13. *A Sturmian walk on Γ^M satisfies the Ferenczi-Zamboni conditions and thus is the symbolic trajectory of some interval exchange with combinatorial data (π_0, π_1^M).*

Proof. Let $(\mu_n, \sigma_n) \in (\Lambda \times \Sigma)^{\mathbb{Z}}$ denote the walk. We simply verify that the six conditions of Theorem 1 are satisfied. Condition (1) is ensured by Lemma 10. Condition (2) is given by Lemma 12. Conditions (3) and (4) are given by Lemma 9.

For condition (5), we remark that since σ is a Sturmian sequence it represents the symbolic trajectory of a point under an irrational rotation and so must satisfy the Ferenczi-Zamboni conditions. In particular, condition (5) holds on the factor σ in our sequence since this is a Sturmian sequence: σ is the symbolic trajectory of an interval exchange with intervals labeled by $\Sigma = \{L, R\}$ with combinatorial data

$$\pi_0(L) = 1 \qquad \pi_1(L) = 2$$
$$\pi_0(R) = 2 \qquad \pi_1(R) = 1$$

and length data $\ell_L = 1 - \theta$, $\ell_R = \theta$ where θ is the rotation parameter described earlier in Section 2.3.

Now suppose that

$$w = (\nu_0, s_0)\,(\nu_1, s_1)\,\cdots\,(\nu_k, s_k)$$

is some finite subword of (μ, σ). Let s denote the Sturmian word $s_0 s_1 \cdots s_k$. The admissible prefixes and suffixes of w are then completely determined by the admissible prefixes and suffixes of s. In particular, there are only two points of concern in showing condition (5) holds, and both of these concerns are alleviated in light of this observation.

1. Suppose $A(w)$ is a singleton, say (ξ, t) is the unique element of $A(w)$. We must show that $D((\xi, t)\, w)$ is a singleton as well. However, if $A((\xi, t)\, w)$ is a singleton, then so is $A(t\, s)$ and since the Ferenczi-Zamboni conditions hold for the sequence (s_n), $D(t\, s)$ must be a singleton. This then implies that $D((\xi, t)\, w)$ is a singleton, since the second component must be the only element in $D(ts)$ and the first one is determined by $\mu_{n+1} = \sigma_n\, \mu_n$. (Note here that the inequalities in condition (5) of the Ferenczi-Zamboni conditions occur simultaneously.)

2. Suppose now that

$$A(w) = \left\{(R^{-1}(\nu_0), R),\, (L^{-1}(\nu_0), L)\right\},$$

where $R^{-1} = h^{-(M+1)}v^{-1}$ and $L^{-1} = h^{-M}v^{-1}$. Then $A(s) = \{R, L\}$. As $(R^{-1}(\nu_0), R) \leq_{\pi_1} (L^{-1}(\nu_0), L)$ and $(s_k \cdot \nu_k, L) \leq_{\pi_0} (s_k \cdot \nu_k, R)$, we must show that if $(s_k \cdot \nu_k, R) \in D((R^{-1}(\nu_0), R)\, w)$, then $D((L^{-1}(\nu_0), L)\, w) = \{(s_k \cdot \nu_k, R)\}$. However, since the Sturmian sequence $(s_n)_{n\in\mathbb{Z}}$ satisfies the Ferenczi-Zamboni conditions, this follows immediately from the fact that if $A(s) = \{L, R\}$ and $R \in D(R\, s)$ then $D(L\, s) = \{R\}$.

To show condition (6), we again remark that the Sturmian sequence σ satisfies the Ferenczi-Zamboni conditions. In particular, condition (6) in the case of σ implies that if w is a finite subword of σ, and $s, s' \in A(w)$ are distinct symbols, then $D(sw) \cap D(s'w)$ must be a singleton. This trivially implies the same for the sequence (μ, σ). □

We are now able to combine the results above to establish our main theorem, restated below for the convenience of the reader.

Theorem 14 (The Characterization Theorem)**.** *Let X be a square-tiled surface on d squares determined by a pair of permutations h and v on $\Lambda = \{1, 2, ..., d\}$. Let $E = \{H, V\}$ be symbols for the edges of the unit square torus. Label the left-hand edge of square λ as (λ, H) and the bottom edge as (λ, V). Then a biinfinite sequence $(\lambda_n, \varepsilon_n) \in (\Lambda \times E)^{\mathbb{Z}}$ is the cutting sequence of an infinite geodesic on X if and only if the following conditions are satisfied:*

(1) $(\lambda_n, \varepsilon_n)$ is consistent with the gluings of the surface;
(2) $(\lambda_n, \varepsilon_n)$ is either periodic, or is minimal but not almost symmetric around a bad square of the surface; and
(3) ε_n is the cutting sequence of a geodesic on the square torus.

The strategy of the proof will be to take a sequence satisfying the three conditions of Theorem 14 and associate to that sequence the symbolic trajectory of an interval exchange on the surface. Such an association necessarily loses some information as two different geodesics with different cutting sequences could give rise to the same symbolic trajectory if those geodesics are related by a Dehn twist. The correspondence between cutting sequences and symbolic trajectories is not one-to-one, but instead depends on a choice which essentially keeps track of how many times Dehn twists were applied.

Proof of Theorem 14. One direction of the proof is simple: it is clear that a cutting sequence $(\lambda_n, \varepsilon_n)$ of a geodesic γ on the surface satisfies the three conditions since the geodesic on X projects to a geodesic on the square torus $\mathbb{T}^2$. By the Veech dichotomy, the geodesic flow in any given direction on X or $\mathbb{T}^2$ is either periodic (if the flow is in a rational direction), or uniquely ergodic (if the flow is in an irrational direction). As Lebesgue measure is preserved by flows in any direction, if the flow is uniquely ergodic the flow must in fact be minimal as it will enter any open ball on the torus. Since the flow on any translation surface is the suspension over an interval exchange (coming from the first-return map of the geodesic flow to an appropriately chosen transversal geodesic segment), minimality of the flow on the surface implies minimality of the interval exchange which implies minimality

of symbolic trajectories. Thus non-periodic cutting sequences must come from geodesics in irrational directions and so are required to be minimal.

For the converse, suppose from $(\lambda_n, \varepsilon_n)$ we construct sequences (μ_n, σ_n) as above. Then by Proposition 13, (μ_n, σ_n) is the symbolic trajectory of some interval exchange with the combinatorial data (π_0, π_1^M) as determined by the surface. We need to show now that the interval exchange on the bases of the squares, as described in Section 3, is one of the interval exchanges for which (μ_n, σ_n) is a symbolic trajectory. By the proof of [FZ08, Theorem 2], this means we need to exhibit a shift-invariant measure on the dynamical system described by (μ_n, σ_n) where the measure on each cylinder set agrees with the lengths of the subintervals of the interval exchange on the surface.

We consider the interval exchange $T_m : I \times \Lambda \to I \times \Lambda$ corresponding to the flow with slope m being the slope of the Sturmian sequence ε as described in Section 2.3.

Define a measure μ on the associated shift space by declaring the measure of a cylinder set to be the Lebesgue measure of the corresponding subinterval in our interval exchange. Denoting the length of an interval J by $|J|$, we are thus considering the measure given by $\mu([w]) = |I_w^0|$. This clearly defines a measure, but we must check that the measure is shift invariant.

Let Δ denote the shift, and let T denote the interval exchange on the surface. We need to show that for each cylinder $[w]$, $\mu([w]) = \mu(\Delta^{-1}([w]))$. Notice

$$\Delta^{-1}([w]) = \bigcup_{a \in A(w)} [aw]$$

where again $A(w)$ is the set of allowable one-character prefixes of the word w. We must show the length $\left|I_w^0\right|$ of the subinterval I_w^0 equals the measure of the Δ-preimage of w.

We may express the measure of the preimage of w as

$$\begin{aligned}
\mu\left(\bigcup_{a \in A(w)} [aw]\right) &= \sum_{a \in A(w)} \left|I_{aw}^0\right| \\
&= \left|\bigcup_{a \in A(w)} \left[I_a^0 \cap T^{-1}(I_w^0)\right]\right| \\
&= \left|\left[\bigcup_{a \in A(w)} I_a^0\right] \cap T^{-1}(I_w^0)\right| \\
&= \left|T\left(\left[\bigcup_{a \in A(w)} I_a^0\right] \cap T^{-1}(I_w^0)\right)\right| \\
&= \left|\left[\bigcup_{a \in A(w)} I_a^1\right] \cap I_w^0\right|.
\end{aligned}$$

By the assumption that the Ferenczi-Zamboni conditions are satisfied, $A(w)$ is a π_1^M-interval, and so $\bigcup_{a \in A(w)} I_a^1$ is a subinterval of I. Since elements of $A(w)$ are precisely the a such that $I_{aw}^0 \neq \emptyset$, the interval $\bigcup_{a \in A(w)} I_a^1$ contains I_w^0 which

implies shift-invariance of our measure μ. Thus T_m is one of the interval exchanges for which the given sequence is a symbolic trajectory.

For each choice of $M \in \{0, 1, 2, ..., D\}$, with D the least common multiple of the lengths of cycles of the permutation h, and each choice of $k \in \mathbb{N}_0$, there is a one-to-one correspondence between symbolic trajectories of the interval exchange on our surface with combinatorial data (π_0, π_1^M) and cutting sequences $(\lambda_n, \varepsilon_n)$ where the length modulo D of the Sturmian sequence (ε_n) is M.

To see this, we simply replace the characters in the sequence as follows: (μ, L) is replaced by

$$(\mu, V)\,(h(\mu), H)\,(h^2(\mu), H)\,\cdots\,(h^{kD+M}(\mu), H),$$

and (μ, R) is replaced by

$$(\mu, V)\,(h(\mu), H)\,(h^2(\mu), H)\,\cdots\,(h^{kD+M}(\mu), H)\,(h^{kD+M+1}(\mu), H).$$

The values of M and k have a very geometric interpretation. Given any geodesic γ that crosses the base of a square λ on the surface, consider the geodesic δ which is obtained by applying a multiple Dehn twist to the surface so that each horizontal cylinder formed by the squares tiling the surface (i.e., corresponding to each cycle of the permutation h) has its top and bottom edges fixed. The geodesics γ and δ certainly have different cutting sequences since δ will cross more vertical edges labeled H than γ will. However, both geodesics have the same symbolic trajectory under the interval exchange on the bases of the squares. See Figure 12 for an example.

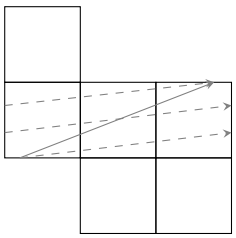

FIGURE 12. Though γ (solid) and δ (dashed) give the same symbolic trajectory of the interval exchange on the surface, they have different cutting sequences.

That is, the sequence (μ_n, σ_n) does not contain enough information to determine a cutting sequence on the surface since two different geodesics, related by an appropriate Dehn twist, will have different cutting sequences but the same (μ_n, σ_n) sequence. The k tells us how many times the Dehn twists are applied and are necessary for recovering the cutting sequence from the symbolic trajectory of the interval exchange.

Finally, note that a sequence which is symmetric about a bad square corresponds to a symbolic trajectory of a discontinuity of the interval exchange. □

6. Concluding Remarks

We end by stating one simple consequence of Theorem 14. If the permutations h and v defining the square-tiled surface are such that $h(\lambda) \neq v(\lambda)$ for each $\lambda \in \Lambda$, then each cutting sequence is uniquely determined by the sequence of labels of

squares, $(\lambda_n)_{n\in\mathbb{N}}$, and the labels of edges can be forgotten. For each λ_i, the next symbol λ_{i+1} will be either $v(\lambda_i)$ or $h(\lambda_i)$. Since $v(\lambda_i) \neq h(\lambda_i)$, the symbol λ_{i+1} uniquely determines the edge crossed by the corresponding geodesic.

As Vincent Delecroix remarked to the author, it may be possible that the sequence of labels $(\lambda_n)_{n\in\mathbb{N}}$ uniquely determines a cutting sequence (λ_n, e_n) even if $h(\lambda) = v(\lambda)$ for some λ. On the other hand, this is not always the case, as shown in the following example due to Pat Hooper. Suppose $\Lambda = \{1, 2, ..., d\}$ and let $h = v$ be the permutation $(1\,2\,3\cdots d)$. The sequence of the labels of edges crossed by any geodesic on the surface is the periodic sequence $(h^n(1))_{n\in\mathbb{Z}}$, and so the cutting sequences are not completely determined by the labels of the squares.

References

[Arn02] Pierre Arnoux. Sturmian Sequences. In *Substitutions in dynamics, arithmetics and combinatorics*, volume 1794 of *Lecture Notes in Mathematics*, chapter 6, pages 143–198. Springer-Verlag, Berlin, 2002. Edited by V. Berthé, S. Ferenczi, C. Mauduit and A. Siegel.

[Dav13] Diana Davis. Cutting sequences, regular polygons, and the Veech group. *Geom. Dedicata*, 162:231–261, 2013.

[Dav14] Diana Davis. Cutting sequences on translation surfaces. *New York J. Math.*, 20:399–429, 2014.

[DPU15] D. Davis, I. Pasquinelli, and C. Ulcigrai. Cutting sequences on Bouw-M\"oller surfaces: an S-adic characterization. *ArXiv e-prints*, September 2015, 1509.03905.

[FZ08] Sébastien Ferenczi and Luca Q. Zamboni. Languages of k-interval exchange transformations. *Bull. Lond. Math. Soc.*, 40(4):705–714, 2008.

[GJ00] Eugene Gutkin and Chris Judge. Affine mappings of translation surfaces: geometry and arithmetic. *Duke Math. J.*, 103(2):191–213, 2000.

[Kea75] Michael Keane. Interval exchange transformations. *Math. Z.*, 141:25–31, 1975.

[KMS86] Steven Kerckhoff, Howard Masur, and John Smillie. Ergodicity of billiard flows and quadratic differentials. *Ann. of Math. (2)*, 124(2):293–311, 1986.

[MH38] Marston Morse and Gustav A. Hedlund. Symbolic Dynamics. *Amer. J. Math.*, 60(4):815–866, 1938.

[MT02] Howard Masur and Serge Tabachnikov. Rational billiards and flat structures. In *Handbook of dynamical systems, Vol. 1A*, pages 1015–1089. North-Holland, Amsterdam, 2002.

[Sch04] Gabriela Schmithüsen. An algorithm for finding the Veech group of an origami. *Experiment. Math.*, 13(4):459–472, 2004.

[Ser85] Caroline Series. The geometry of Markoff numbers. *Math. Intelligencer*, 7(3):20–29, 1985.

[SU11] John Smillie and Corinna Ulcigrai. Beyond Sturmian sequences: coding linear trajectories in the regular octagon. *Proc. Lond. Math. Soc. (3)*, 102(2):291–340, 2011.

[Vee89] W. A. Veech. Teichmüller curves in moduli space, Eisenstein series and an application to triangular billiards. *Invent. Math.*, 97(3):553–583, 1989.

[Wri15] Alex Wright. Translation surfaces and their orbit closures: an introduction for a broad audience. *EMS Surv. Math. Sci.*, 2(1):63–108, 2015.

[WZ15] ShengJian Wu and YuMin Zhong. On cutting sequences of the L-shaped translation surface tiled by three squares. *Sci. China Math.*, 58(6):1311–1326, 2015.

[Zor06] Anton Zorich. Flat surfaces. In *Frontiers in number theory, physics, and geometry. I*, pages 437–583. Springer, Berlin, 2006.